The Rise of People-Centric Sensing

A. T. Campbell*, S. B. Eisenman†, N. D. Lane*, E. Miluzzo*, R. A. Peterson*
H. Lu*, X. Zheng*, M. Musolesi*, K. Fodor*, G.-S. Ahn†

†Columbia University, New York, NY, USA.
*Dartmouth College, Hanover, NH, USA.

Abstract

People-centric sensing is poised to radically change the way we see the world. Technological advances in sensing, computation, storage, and communications will turn the near ubiquitous mobile phone into a global mobile sensing device that enables myriad new personal, social, and public sensing applications. People-centric sensing will help drive this trend by enabling a different way to sense, learn, visualize and share information about ourselves, friends, communities, the way we live and the world we live in. People-centric sensing juxtaposes the traditional view of mesh sensor networks where people are passive data consumers that simply interact at the network periphery with physically embedded static sensor webs, with one where people carry mobile sensing elements (i.e., sensor-enabled mobile phones), enabling opportunistic *sensing coverage, and thus represent a key architectural component of the system. In this article, we discuss our vision for people-centric sensing, the challenges it brings, and the ongoing development of a number of social sensing applications as part of the MetroSense Project.*

1 Introduction

The evolution of sensing, computing and communication technology over the past few years has brought us to a tipping point in the field of wireless sensor networking. A decade ago, research prototype hardware began to emerge facilitating the genesis of wireless sensor networks as they exist today - small resource-limited embedded devices that communicate via low-power low-bandwidth radio. A natural first application of these networks of custom devices was solving relatively small-scale specialized problems in the scientific and industrial domains, such as forest monitoring and preventative maintenance. While these problems and applications remain important, the recent miniaturization and subsequent introduction of sensors into popular consumer electronics like mobile phones (e.g., Apple iPhone), PDAs (e.g., Nokia N810), and mp3 players (e.g., Nike+iPod) has opened the door to a new world of application possibilities. With wireless sensor platforms in the hands of the masses, and with the proper architectural support, wireless sensor networks can be leveraged to address urban-scale problems or provide global information access (i.e., public sensing applications). At the same time, people as individuals, or in social or special interest groups, can apply these new sensing networks to applications with a more personal focus. We see a continuing push in this direction and the advent of a new era of people-centric sensing.

In a people-centric sensing system, humans, rather than trees or machines, are the focal point of sensing, and the visualization of sensor-based information is for the benefit of common citizens and their friends, rather than domain scientists or plant engineers. Additionally, it is the aggregate mobility of the humans themselves that enables both sensing coverage of large public spaces over time, and allows an individual, as a custodian of the sensing device, to collect very targeted information about his or her daily life patterns and interactions. We say the sensing coverage of spaces, events, and human interactions is *opportunistic* in a people-centric system since the system architecture has no point of control over the human mobility patterns and actions that facilitate this coverage. While this lack of control can translate to gaps in sensing coverage, the alternative of a world-wide web of static sensors is clearly not tenable in terms of monetary cost, scalability, or management. Further, by having the sensing devices carried by people (e.g., mobile phones, iPods, etc.), a people-centric sensing system creates a symbiotic relationship between itself and the communities and individuals it serves.

In this article, we describe our vision of people-centric sensing, and the architectural support we are developing in the MetroSense Project [2] to realize this vision. People-centric sensing (see Figure 1) gives rise to a host of new applications that can be classified into three main groups: (i) *personal sensing*, those focused on personal monitoring and archiving; (ii) *social sensing*, those where information is shared within social and special interest groups; and (iii) *public sensing*, those where data is shared with everyone for the greater public good (e.g., entertainment, community action). Each of these application foci comes with its own challenges in terms of how data is best sampled, understood (e.g., via mining, inference), visualized and shared with others. We present a number of prototype applications we are developing in the MetroSense Project that cover these sensing scenarios, including a discussion of how people can best learn information of interest from the raw data, represent that information in a meaningful way and share that information, as appropriate.

2 Background

People-centric sensing sits at the nexus of several research disciplines, including sensor networking, pervasive computing, mobile computing, machine learning, human-computer interfacing, and social networking. Significant research contributions made within each discipline have facilitated the rise of people-centric sensing, and research focusing on synthesizing these contributions is now emerging [1]. In the following, we highlight current projects related to our people-centric sensing initiative. SensorPlanet [15] is a Nokia-initiated global research framework for mobile-device-centric wireless sensor networks. SensorPlanet provides hardware platforms and a research environment that enable the collection of sensor data on a large and heterogeneous scale, and establishes a central repository for sharing the collected sensor data. SenseWeb [5], a project sponsored by Microsoft Research, provides shared sensing resources and sensor querying and data collection mechanisms to develop sensing applications. In both projects, participating universities develop their own applications and share the collected data to facilitate research on data analysis and mining, visualization, and machine learning. The Urban UCLA Sensing initiative [13] has a vision of equipping users to compose a sensor-based recording of their experiences and environment by leveraging sensors embedded in mobile devices and integrating existing public outlets of urban information (e.g., weather, traffic, air quality). Urban Sensing is exploring how these individual stories of everyday life can be coordinated to document the urban environment, as well as be fused with other sensed data about the city and fed back into the physical, collective experience in urban public spaces. The Intel-sponsored Urban Atmospheres project [14] is also using sensors to explore the human condition. The MIT Cartel project [4] provides a mobile communications infrastructure based on car-mounted communication platforms exploiting open WiFi access points in a city, and provides urban sensing information such as traffic conditions. The Harvard/BBN CitySense project [11] provides a static sensor mesh offering similar types of urban sensing data feeds.

3 The MetroSense Vision

The MetroSense conception of a people-centric sensing system is based on a three stage Sense, Learn, Share framework. In the sense stage, MetroSense leverages mobility-enabled interactions between human-carried mobile sensors (e.g., mobile phones, personal medical sensing devices), static sensors embedded in the civic infrastructure (e.g., vehicle-based sensing networks, home medical sensing networks), and edge wireless access nodes providing a gateway to the Internet, to support: (i) the delivery of application requests to the mobile devices, (ii) the sampling of sensors specified by the request, and (iii) the delivery of sampled data back to the application. Application functions, including generating requests, data

Figure 1. People-centric sensing applications can be thought of as having a personal, social, or public focus.

analysis and visualization logic, may be installed directly on the mobile device, may run on remote servers (e.g., a web application) but communicate with the mobile device via wireless gateway nodes (e.g., GPRS gateway, WiFi access point), or may be split between the mobile device and these servers. An application sampling request specifies at least one required sensor type (e.g., an accelerometer), and the required sampling context, that is, the set of conditions that must be met for the sampling to take place (e.g., time of day, location, sensor orientation). In the learn stage, we analyze the sensed data using both simple statistical measures and more involved machine learning techniques to extract higher-level meaning. The choices of what data analysis techniques to apply, and what data features to analyze are made to best match the availability and characteristics of the sensed data (e.g., noisiness, incompleteness), and the target application visualization. In addition, we leverage the people-centric nature of the system by using social connections between system users when possible to improve the performance of learned models (e.g., activity classifiers) and decrease the time it takes to learn these models. In the share stage, the learned information is visualized by the individual and optionally shared. For example, sharing is possible within social groups (e.g., classroom, sports team, Facebook, MySpace), or within a global community (e.g., Second Life). Each of the stages in the Sense, Learn, Share framework are further discussed in the following.

4 Sense: Exploit Mobility, Be Opportunistic

From an infrastructural point of view, we observe that sensing, computing, and communication resources are already widely deployed, in the form of end user electronics, enterprise and public radio access networks, and Internet backhaul. Therefore, there is no longer a need to deploy specialized custom-built hardware to enable collection and transport of people-centric sensor data. Instead, we aim to symbiotically leverage this extant infrastructure to support people-centric applications. We also take advantage of the increasing integration of sensors (e.g., camera, microphone, accelerometer, temperature, infrared) into off-the-shelf consumer devices (e.g., Nokia N95 mobile phone, Nokia N810 tablet PCs, iPod Touch) to transparently sample the device custodian (i.e., the human user) and the custodian's environment.

4.1 What Role do People Play?

Utilizing human-carried devices as a fundamental building block of the sensing system raises the question of what roles people, as sensing device custodians, should (or are willing to) play in the architecture. Deciding to what extent people should be conscious active participants in meeting application sensing requirements has significant design implications, especially in terms of defining the fundamental research challenges to implementing a robust, scalable and secure system. It is useful to consider the two endpoints of the design spectrum of custodian awareness and involvement [8], which can be thought of as opportunistic [2] and participatory [13].

With participatory sensing the custodian consciously opts to meet an application request out of personal or financial interest. A participatory approach incorporates people into significant decision stages of the sensing system, actively deciding what application requests to accept, what data to share and to what extent privacy mechanisms should be allowed to impact data fidelity. Since most of the tough sensing and privacy decisions are made by the human, a purely participatory system design focuses on tools that assist people to share, publish, search, interpret and verify information collected using a custodian's device. Purely participatory sensing places many demands on involved device custodians (e.g., prompting via their device GUI for authorization to take a sound sample or share a particular image sample), which restricts the pool of willing participants. The tolerance of people to endure interruptions on behalf of applications limits the number and request load of concurrent applications that can likely be supported. Further, under the participatory approach, an application needs to have a critical mass of community appeal. These factors may combine to limit both an application's scale and the diversity of applications that are likely to be supported by a purely participatory people-centric network.

In the MetroSense Project, we emphasize an *opportunistic approach*. Opportunistic sensing shifts the burden of supporting an application from the custodian to the sensing system, automatically determining when devices can be used to meet application requests. In this paradigm, the custodian configures his device to allow applications to run (subject to privacy and resource usage restrictions), but may not be aware of the applications active at any given time. Instead, a custodian's device (e.g., mobile phone) is utilized whenever its state (e.g., geographic location, body location) matches the context requirements of an application. In this way, applications can leverage the sensing capabilities of all system users without requiring human intervention to actively and consciously participate in the application, lowering the bar for applications to run in people-centric networks. To support symbiosis between the custodian and the system, sensor sampling occurs only if the privacy and transparency needs of the custodian are met. The main privacy concern is the potential leak of personally sensitive information indirectly when providing sensor data (e.g., the custodian's location). To maintain transparency, opportunistic use of a device should not noticeably impact the normal user experience of the device.

4.2 Challenges of Opportunistic Sensing

Along with the aforementioned benefits, the opportunistic paradigm introduces a number of challenges. The opportunistic use of sensor custodian devices means that sensing is a secondary, low-priority operation on the mobile device (e.g., mobile phone). Consequently, the periods when the device is able to meet the sensing requirements defined by an application request may be short and intermittent. Opportunistic systems also take on much more of the decision making responsibility and are thus more complex and may use more resources. Specific challenges that must be overcome for opportunistic sensing to be feasible include: determining the sampling context of the device; adapting to the changing resource availability and sampling context of devices; achieving sufficient sensing coverage in the face of sensing target mobility; and sustaining custodian privacy. We are currently investigating methods to address these challenges, as outlined in the following.

4.2.1 Sensing Context

Sensing context is the meta-data that describes the conditions to which the sensing hardware is exposed and affects both the sensor data itself and its ability to perform the sensing operation. Knowledge of sensing context is required as an input to a number of operations of an opportunistic sensing system. With sensor sharing (described later), it is used to evaluate potential candidate sensor devices in terms of a given application request. During servicing of application requests it indicates when sampling should be started and stopped. More generally, knowing sensing context is important in understanding the sampled data, especially from consumer devices (e.g., mobile phone) where sensing is largely a second class citizen and samples may be taken under suboptimal conditions (e.g., when the device is in a pocket or purse).

4.2.2 Sensor Sharing

In order for an opportunistic sensing system to collect samples that meet a general set of application requirements (e.g., sensor type, location, physical orientation, time), it must be able to adapt to the changing resource availability and sampling context of sensing devices. For example, a mobile phone device may run out of memory or power, or may be placed in the custodian's pocket before a required light level, sound, GPS, or image sample is taken. To help the sensing system be more robust to these changes, we are developing a sensor sharing mechanism. This approach allows application requests assigned to a particular device to borrow samples from the best-suited sensors (i.e., matching the required sensing context, not already in use by another application on the device) of any available device in the neighborhood at the time. Devices exchange current context information, and data is selected from the device whose context most closely matches the requirements of the application. Given the potentially rapid dynamism of sampling context, a research challenge is determining a context matching metric that, when used for this sharing mechanism, provides samples with high average-case fidelity with respect to the applications requirements.

4.2.3 Mobile Target Sensing

In people-centric sensing where people are frequently the intended sensing targets, there is a need to support the tracking and sensing of mobile targets (e.g., a noisy truck, a missing child's voice) with mobile sensing devices. There are two major challenges in building mobile event sensing system using mobile sensors carried by people. First, mobile sensors need to be informed about the sensing target (i.e., be "tasked") before sensing, but for efficiency only those mobile sensors near the mobile target should be tasked. Second, there is no guarantee that there will always be enough mobile sensors around the target to maintain sensing coverage. To efficiently establish a sensing area around the target, a mobile sensor that detects the target using its sensors forwards the task to its neighbors. To recover a lost target, we estimate the area to which the target is predicted to move based on a distributed Kalman filter, and then use a geocast scheme to forward the task to the sensors in the predicted area.

4.2.4 Privacy

Opportunistic sensing faces barriers to wide scale adoption unless users trust the system to provide privacy guarantees on par with those provided by state of the art systems, a difficult research challenge. For the sensing device custodians, the potential exists for the leakage of sensitive personal information both from the collected data samples themselves, and from the process by which the samples are collected. As an example of the latter, during sensor sharing, the fact that data is shared between devices reveals information about the contexts of these devices. Also, sensor data itself (e.g., images, sound, accelerometer data) may contain information that device custodians do not wish to expose about themselves. Further, even those who are not device custodians and may not be the primary sensing targets are vulnerable to an accidental compromise of privacy, a "second-hand smoke" of people-centric sensing systems. Ongoing work in the MetroSense project begins to address these issues by providing mobile device custodians with a notion of anonymity through k-anonymous tasking [6].

5 Learn: Understanding Opportunistic Data

Once the data has been sampled and delivered to the application, higher level meaning must be extracted from the raw samples. Two main challenges facing data analysis in the people-centric sensing domain [9], and in the construction of accurate inference models (e.g., human activity models) in general, are the lack of appropriate sensor data inputs and the time and effort that must be spent in training models that gives sufficient classification accuracy. The COTS-device-based sensing substrate upon which we build people-centric applications is characterized by heterogeneity in terms of sensing and other resources (e.g., memory, battery capacity, CPU power), impacting both the construction and usage of models. The data inputs most useful in generating high accuracy models may not be available on all devices, requiring users of less capable devices to settle for less accurate models based on other available data features. As an example using a snapshot of current technology, a common approach in the literature is to extract data features from a GPS sensor to generate an indoor/outdoor classifier. However, GPS is integrated into only a relatively small percentage of mobile phones on the US market today. These observations motivate a careful consideration of where in the architecture classifiers should run, and inspire two possible solutions for model creation: opportunistic feature vector merging, and social-network-driven sharing of models and training data.

5.1 Opportunistic Feature Vector Merging

With the opportunistic feature vector merging approach we seek to push the performance of classification models possible with sensor-poor devices towards that possible with sensor-rich devices. When merging feature vectors (i.e., multi-element numerical object or activity representations), data features available from more capable devices are borrowed and merged with data features natively available from a less capable device in the model building stage, allowing a higher accuracy model to be built even for the less capable device. This borrowing is facilitated by opportunistic interaction, both direct and indirect, between a less capable device and a more capable device in situ. As an example of direct interaction, as two mobile phone users follow their daily routines, a mobile phone without GPS can borrow GPS data features from a mobile phone with GPS as an input to its indoor/outdoor location classifier. For indirect interaction, both devices collect data samples according to their respective capabilities. Subsequently, centralized matching between other features collected by both devices (i.e., non-GPS-based features) may provide for a binding between the feature vector collected by the phone without GPS and the GPS features collected by the GPS-equipped phone. The GPS features can then essentially be borrowed via this binding.

5.2 Social-network-driven Model and Data Sharing

Even when devices provide an appropriate set of data features to build accurate models, users may be required to gather a large set of training data (perhaps manually labeling it) before applications using the models' outputs work best. The inconvenience in both the labeling of training data and the time required for model training to complete may act as disincentives to the broad-scale adoption of new people-centric applications. We propose the sharing of training data among users to reduce training time and labeling effort by amortizing the model training cost over all system users. However, this is likely to reduce the accuracy of the resulting model (e.g., since people do the same activity in many slightly different ways, and might describe the same activity with slightly different labels). With our social-network-driven sharing approach, training data is shared only within social circles, within which we conjecture group vocabularies and other commonalities lead to more consistent and understandable labeled training data and a higher model accuracy, while still reducing the quantity of per user training data required. A careful consideration of the particular labeling problem is required in deciding within which social group sharing might be most effective. Initial results implementing these two techniques are promising [7].

5.3 Additional Resource Considerations

In addition to model generation, resource limitations on mobile devices designed primarily for other purposes require a careful consideration of where data processing takes place. For example, due to limitations in CPU power we have noticed running a full spectrum FFT on a mobile phone can impact other ongoing operations, and can run too slowly to keep to keep up with the stream of sampled data. Such behaviour violates our tenet of symbiosis with the device's primary user experience. Further, due to local storage limitations, analysis that requires access to a large amount of historical data may not be possible without interaction with persistent storage on backend servers. When placing learning functionality, a systemic view is required that considers mobile phone resource constraints, communication cost to the backend servers, and the sampling rate required to detect and characterize the phenomena of interest.

6 Share: Enabling Social Sensing Applications

People-centric sensor networking aims to support applications that engage the general public. This potential for interest across a broader segment of the population, in concert with our use of an opportunistic sensing design, facilitates the availability of a massive number of mobile sensing devices, in turn increasing the scope and scale of applications that can be supported, and also improving the fidelity of sampled objects, events, and human activities. In the following, we discuss a number of applications we have developed in the MetroSense Project that incorporate personal, social and public sensing.

6.1 CenceMe - Social Sensing

The growing ubiquity of the Internet provides the opportunity for an unprecedented exchange of information on a global scale. Those with access to this communication substrate, and among these especially the youth, increasingly incorporate personal information exchange (e.g., availability, activity, mood) into their daily routines via technologies such as email, blog, instant message (e.g., Skype or MSN), SMS, video sharing (e.g., YouTube), social network software (e.g., Facebook and MySpace), and VOIP. Yet, the question of how to incorporate personal sensing information such as human activity inferencing into these applications has remained largely unexplored. While existing communication forums allow the exchange of text, photos, and video clips, we believe a more richly textured user experience can be provided by integrating *automatically* harvested, processed and shared sensor data into the mix. With the CenceMe application [10], we distill this sensed data (see Figure 2(a)) into what we call a user's "sensing presence", a virtual representation of a user's status in terms of his activity (e.g., sitting, walking, meeting friends), disposition (e.g., happy, sad, ok), habits (e.g., at the gym, coffee shop, at work) and surroundings (e.g., noisy, hot, bright, high ozone).

We are evolving a prototype implementation of CenceMe that enables members of social networks both to access historical traces of their own data, and more powerfully to share their sensing presence among their buddies in a secure manner. For users on-the-go, we have implemented a client to show current buddy sensing presence on the GUI-based displays of most new mobile phones, using a set of simple and intuitive icons representing, for example, the activity and location of a user. We include this sensing presence snapshot, along with a more complete, browsable representation of the sensing presence of a user and his buddies via the CenceMe web portal. Examples include archived historical traces, comparisons of presence attributes with friends - am I above average?, and extracted patterns and features of importance in one's life routine, etc. However, the real power of the automatic inference of sensing presence is the ability for a user to configure CenceMe to export this sensing presence, without direct intervention, across his online social networks. This has only recently been made possible through the release of developer APIs for Facebook, Skype, Pidgin, MSN Messenger and the like. We have implemented a number of widgets a user can add to her Facebook account (see Figure 2(b)) to share various representations of sensing presence with her Facebook buddies. CenceMe users share data according to CenceMe group membership policies set through the web portal. CenceMe buddies are defined by the combination of buddy lists imported from the social networking application accounts a user registers with CenceMe. Thus, CenceMe inherits group structures already created by a user for his other applications, but also allows for the specification of more so-

(a) CenceMe distills a user's "sensing presence" from samples taken from sensors embedded in personal mobile devices (e.g., mobile phone, PDA, mp3 player), sports equipment (e.g., running shoes, bicycle), and the civic infrastructure. Sensing presence is sharable with a user's friends through popular social networking applications.

(b) We have built widgets for Facebook that allow expression of sensing presence through the friends list, the mini-feed, and a dedicated Sensor Presence display.

Figure 2.

phisticated group privacy policies within CenceMe.

6.2 Global Sharing in Virtual Worlds

Virtual world simulators (e.g., Second Life) represent one of the new frontiers in online entertainment and business services. People lead virtual lives in these alternative worlds using personal avatars. Bridging real life and these virtual worlds together is challenging, but enables new application scenarios for these systems via public sensing. The sensors embedded in commercial mobile phones can be used to infer real-world activities, that in turn can be reproduced (see Figure 3) in public virtual environments, sharing with the (virtual) world. While previous research has brought static objects from the real world to the virtual world, we are the first to bring people's real world sensing presence to active subjects (i.e., their avatars) in the virtual world [12]. People's sensor data can be rendered in the virtual world anywhere on the spectrum between reality and fantasy. That is, arbitrary mappings between sensed physical data and the avatar actions, appearance, and location are possible. Further, the connection between the physical and virtual worlds need not be only one way, and we envision that users may receive communication (e.g., emails, instant messages or SMS) or actuation triggers (e.g., mobile phone vibration) to indicate the status or environment changes experienced by their avatar in the virtual world. As an initial step, we have implemented activity recognition and voice detection classifiers that run on the mobile phone, acting on data from local embedded sensors. We have also built a data bridge using available APIs to control a user's avatar in Second Life (see Figure 3).

6.3 BikeNet - Recreational Sensing

BikeNet is a recreational application that contains elements of personal, social and public sensing. There is substantial interest in the mainstream recreational cycling community in collecting data quantifying various aspects of the cycling experience, mirroring the broader interest in fitness metrics among exercise enthusiasts and other health conscious individuals. Existing commercial bike-sensing systems targeting this

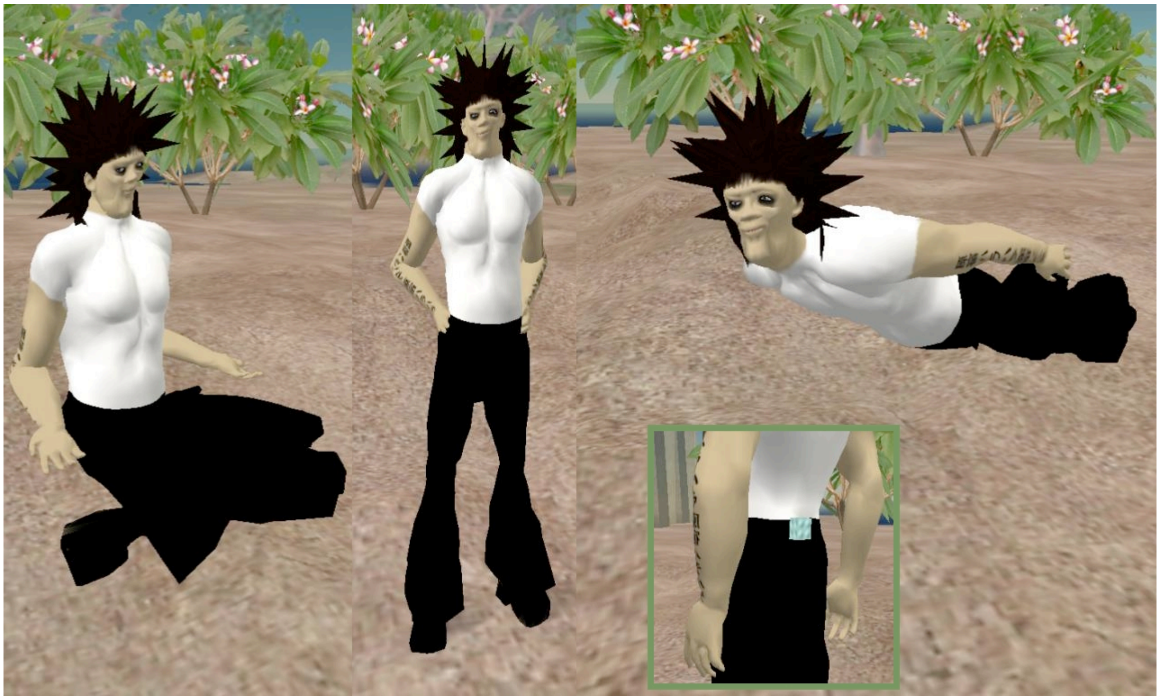

Figure 3. Second Life integration with the physical world. Accelerometer data is collected from a person's mobile phone, and classified into activity states sitting, standing and running. These states are then injected into Second Life via the mobile phone object the avatar carries (see picture inset). The Second Life user defines the profile for his avatar to interpret and render these incoming activity states. For example, in the figure sitting, standing and running have been mapped by the user to yoga-floating, standing, and flying, respectively.

demographic measure and display simple data such as wheel speed, and provide simple inferences such as distance traveled and calories burned. These systems have become increasingly more sophisticated and miniaturized.

We have designed and implemented the prototype of a system reflecting a future where wirelessly accessible sensors are commonly embedded in commercially manufactured bicycles, and the cyclist's mobile device (e.g., mobile phone) interacts with these sensors during the ride to quantify aspects of cycling performance and environmental conditions. In terms of personal sensing, we view this system as akin to the Nike+iPod kit, a system for recreational runners that logs exercise history. The BikeNet application [3] measures the following metrics to give a holistic picture of the cyclist experience: current speed, average speed, distance traveled, calories burned, path incline, heart rate, CO_2 level, car density surrounding the cyclist, and galvanic skin response (a simple indicator of emotional excitement or stress level). All data sensed by the system is stamped with time and location metadata. This data is provided to the cyclist immediately, for example, via the mobile phone's LCD, but is also uploaded to a personal repository on remote BikeNet servers for long term archiving, and for later trend analysis (e.g., cycling performance, personal health). The BikeView portal (http://bikenet.cs.dartmouth.edu) provides a personal sensing repository for all ride statistics. Additionally, it provides for sharing of information via real-time requests, a function most likely enabled within social groups (e.g., for race-time management within cycling team, or for rough tracking of a family member to know when to order the pizza). Sharing of aggregate statistics and route rankings (optionally stripped of identifying information) is facil-

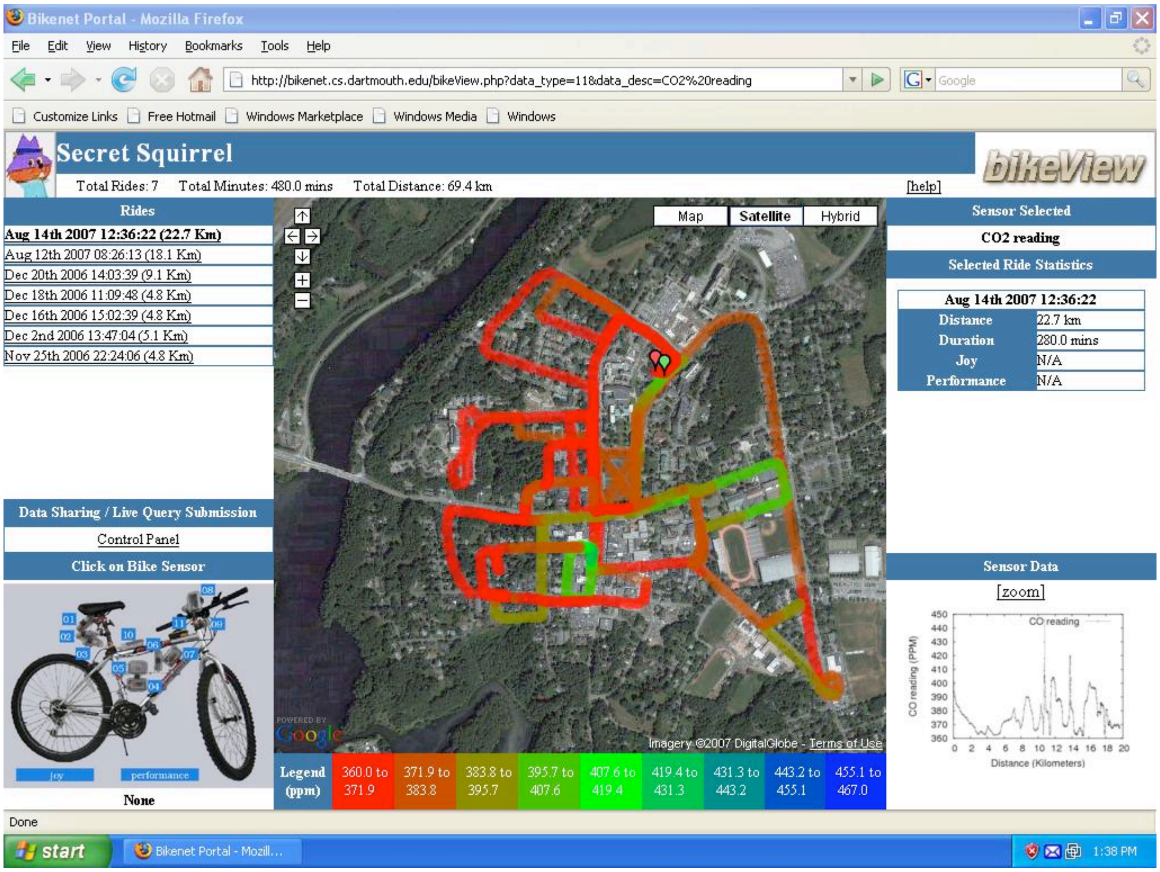

Figure 4. The BikeView portal provides personal access to archived cycling data, which can be socially shared with cyclists, or used to support a public sensing initiative. The CO_2 map shown here is the result of multiple users' data merged to form a complete map of Hanover, NH, USA.

itated within the cycling community group. In addition, BikeNet facilitates public sensing and sharing by allowing multiple users to merge their individual data, for example, to create pollution, allergen, and noise maps of their city. Such a map not only provides a way to learn about the safest (e.g., least cars encountered) and healthiest (e.g., better air quality, lower noise) ways to get around town, but also may provide the basis for political action to improve the city. Figure 4 shows such a map, built through the BikeView portal, of CO_2 readings combined from users of our prototype BikeNet system mapping the town of Hanover, NH, USA.

7 Conclusion

Technology advances in sensing and microelectronics and their integration into everyday consumer devices lays the ground work for the rise of people centric sensing. Applications will be structured around personal, group, and public sensing. The MetroSense Project (http://metrosense.cs.dartmouth.edu) is developing new architectures, protocols, and applications for people-centric sensing, where people are used directly as mobile sensing facilitators and indirectly, by leveraging their social network structures, in the learning and sharing of sensing presence.

Acknowledgment

This work is supported in part by Intel Corp., Nokia, NSF NCS-0631289, the Institute for Security Technology Studies (ISTS) at Dartmouth College, and by members of the MetroSense project. ISTS support is provided by the U.S. Department of Homeland Security under Grant Award Number 2006-CS-001-000001. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security.

References

[1] T. Abdelzaher, Y. Anokwa, P. Boda, J. Burke, D. Estrin, L. Guibas, A. Kansal, S. Madden, J. Reich. Mobiscopes for Human Spaces. *IEEE Pervasive Computing*, Vol. 6, No. 2, pp. 20-29, April-June, 2007.

[2] A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo and R. A. Peterson. People-Centric Urban Sensing (Invited Paper). In *Proc. of 2nd ACM/IEEE Int'l Conf. on Wireless Internet*, Boston, August 2-5, 2006.

[3] S. B. Eisenman, E. Miluzzo, N. D. Lane, R. A. Peterson, G.-S. Ahn, and A. T. Campbell. The BikeNet Mobile Sensing System for Cyclist Experience Mapping. In *Proc. of 5th ACM Conf. on Embedded Networked Sensor Systems*, Sydney, November 6-9, 2007.

[4] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. Miu, E. Shih, H. Balakrishnan, and S. Madden. CarTel: A Distributed Mobile Sensor Computing System. In *Proc. of 4th Int'l Conf. on Embedded Networked Sensor Systems*, pp. 125–138, Boulder, November 1-3, 2006.

[5] A. Kansal, S. Nath, J. Liu and F. Zhao. SenseWeb: An Infrastructure for Shared Sensing. In *IEEE Multimedia*, Vol. 14, No. 4, pp. 8-13, October-December 2007.

[6] A. Kapadia, N. Triandopoulos, C. Cornelius, D. Peebles and D. Kotz. AnonySense: Opportunistic and Privacy-Preserving Context Collection. In *Proc. of 6th Int'l Conf. on Pervasive Computing*, Sydney, May 19-22, 2008.

[7] N. D. Lane, H. Lu, S. B. Eisenman and A. T. Campbell. Cooperative Techniques Supporting Sensor-based People-centric Inferencing. In *Proc. of 6th Int'l Conf. on Pervasive Computing*, Sydney, May 19-22, 2008.

[8] N. D. Lane, S. B. Eisenman, M. Musolesi, E. Miluzzo and A. T. Campbell. Urban Sensing Systems: Opportunistic or Participatory?. In *Proc. of 9th Workshop on Mobile Computing Systems and Applications*, Napa Valley, February 25-26, 2008.

[9] J. Lester, T. Choudhury and G. Borriello. A Practical Approach to Recognizing Physical Activities. In *Proc. of 4th Int'l Conf. on Perv. Comp.*, pp. 1-16, Dublin, May 2006.

[10] E. Miluzzo, N. D. Lane, S. B. Eisenman, and A. T. Campbell. CenceMe - Injecting Sensing Presence into Social Networking Applications. In *Proc. of EuroSSC 2007*, Springer LNCS 4793, pp. 1–28, 2007.

[11] R. Murty, A. Gosain, M. Tierney, A. Brody, A. Fahad, J. Bers, and M. Welsh. CitySense: A Vision for an Urban-Scale Wireless Networking Testbed Harvard University Technical Report TR-13-07, September, 2007.

[12] M. Musolesi, E. Miluzzo, N. D. Lane, S. B. Eisenman, T. Choudhury, A. T. Campbell. The Second Life of a Sensor - Integrating Real-world Experience in Virtual Worlds using Mobile Phones. Technical Report, Dept. of Computer Science, Dartmouth College, October, 2007.

[13] A. Parker, S. Reddy, T. Schmid, K. Chang, G. Saurabh, M. Srivastava, M. Hansen, J. Burke, D. Estrin, M. Allman and V. Paxson. Network System Challenges in Selective Sharing and Verification for Personal, Social, and Urban-scale Sensing Applications. In *Proc. of HotNets-V*, Irvine, November 2006.

[14] E. Paulos and T. Jenkins. Urban Probes: Encountering Our Emerging Urban Atmospheres. In *Proc. of SIGCHI Conf. on Human Factors in Computing Systems*, Portland, April 2-7, 2005.

[15] V. Tuulos, J. Scheible and H. Nyholm. Combining Web, Mobile Phones and Public Displays in Large-Scale: Manhattan Story Mashup. In *Proc. of the 5th Int'l Conf. on Pervasive Computing*, Toronto, May 2007.